

related parameters that may be considered in the operation of such logical management structures include any parameter that at least partially characterizes one or more aspects of a particular memory unit including, but are not limited to, parameters such as recency, frequency, aging time, sitting time, size, fetch (cost), operator-assigned priority keys, status of active connections or requests for a memory unit, etc.

FIG. 1 is a simplified representation of one embodiment of a storage system 100 including a storage management processing engine 105 coupled to storage devices 110 using, for example, fiber channel loop 120 or any other suitable interconnection technology. In the illustrated embodiment, storage devices 110 may include a plurality of storage media, for example, a group of storage disks provided in a JBOD configuration. However, it will be understood that storage devices 110 may include any other type, or combination of types, of storage devices including, but not limited to, magnetic disk, optical disk, laser disk, *etc.* It is also possible that multiple groups of storage devices may be coupled to storage management processing engine 105. Further, it will be understood that although storage system embodiments are illustrated herein, that benefits of the disclosed methods and systems may be realized in any information management system I/O resource environment including, but not limited to, storage system environments, file system environments, *etc.*

As shown in FIG. 1, storage management processing engine 105 may include an I/O manager 140 that may receive requests, *e.g.*, from a file subsystem, for information or data contained in storage devices 110. I/O manager 140 may be provided with access to I/O characteristic information, for example, in the form of an I/O capacity data table 145 that includes such information as the estimated average access delay and the average transfer rate per storage device type, storage device manufacturer, fiber channel topology, block size, *etc.* In one exemplary embodiment, I/O manager 140 and an I/O capacity data table 145 may be combined as part of a storage sub-processor resource manager. Storage management processing engine 105 may also include a cache/buffer manager 130 that monitors or is otherwise aware of available buffer memory in storage system 100.

In the illustrated embodiment, I/O manager 140 may be configured to be capable of monitoring or tracking one or more system I/O performance characteristics of an information management system including, but not limited to, average access delay (“AA”), average transfer rate (“TR”) from the storage device(s) to the I/O controller, total number of viewers (“NoV”) 5 each presenting an I/O request for a continuous portion of an information object such as a multimedia object, consumption rate (“P”) of one or more viewers in the I/O queue, the maximal available buffer memory (“ B_{max} ”), combinations thereof, *etc.*

In one exemplary implementation of the embodiment of FIG. 1, storage devices 110 may 10 be JBOD disks provided for delivering continuous content such as streaming video. In this implementation, I/O manager 140 may be configured to have substantially total control of the I/O resources, for example, using one single JBOD with two fiber channel arbitrated loop (“FC-AL”) loops 120 shared by at least two information management systems, such as content delivery or router systems such that each information management system may be provided with 15 a number of dedicated storage devices 110 to serve its respective I/O workload. However, it will be understood that the disclosed methods and systems may be implemented with a variety of other information management system I/O resource configurations including, but not limited to, with a single information management system, with multiple JBODs, combinations thereof, *etc.* It will also be understood with benefit of this disclosure by those of skill in the art that the 20 configuration of storage devices 110 may be optimized in one or more ways using, for example, disk mirroring, redundant array of independent disks (“RAID”) configuration with no mirroring, “smart configuration” technology, auto duplication of hot spots, and/or any other method of 25 optimizing information allocation among storage devices 110.

In this exemplary implementation, AA may be estimated based on average seek time and 25 rotational delay, TR may be the average transfer rate from storage device 110 to the reading arm and across fiber channel 120, B_{max} may be obtained from cache/buffer manager 130, and P_i may represent the consumption rate for a viewer, *i*, in the I/O queue which may be impacted, for example, by the client bandwidth, the video playback rate, *etc.* For illustration purposes, block 30 size (“BL”) is assumed to be constant for this exemplary implementation, although it will be understood that variable block sizes may also be employed in the practice of the disclosed

methods and systems. Also for purposes of describing this exemplary embodiment, it is assumed that within one control relationship between the illustrated set of storage devices 110 and I/O manager 140, all existing viewers are serviced in a sequence of cycles or rounds. In each such cycle or round it is assumed that each viewer, i , is served once by fetching N_i number of blocks of a given segment size, followed by a waiting period until the next cycle where another N_i number of blocks is fetched to serve each viewer, i . The time interval of each such cycle may be denoted by “ T ”.

In this example, uneven distribution of total I/O demands among a number (“NoD”) of multiple storage devices 110 may be quantified or otherwise described in terms of a factor referred to herein as “Workload Skew” (represented herein as “Skew”), the value of which reflects maximum anticipated retrieval demand allocation for a given storage device 110 expressed relative to an even retrieval demand distribution among the total number of storage devices 110. For example, assuming a NoD value of 10 storage devices 110, and assuming a total NoV of 1000 viewers, an even demand distribution would be 100 viewers (*i.e.*, a Skew value of 1) per storage device 110. A Skew value of about 1 may be encountered, for example, in systems that employ substantially completely mirrored storage disks, RAID storage disks, or that employ other intelligent disk allocation technology.

Skew values may be an actual value that is measured or monitored, or may be a value that is estimated or assumed (*e.g.*, values assumed for design purposes). Skew values of greater than about 1 may be encountered in less optimized system configurations. For example, if two given disk drives out of a 10 disk drive system have a “hot” Skew value (*e.g.*, Skew value of from about 2 to about 4), this means that most new requests are going to the two given disk drives, which are fully utilized. Under these conditions, system optimization is limited to the two heavily loaded disk drives, with the remaining 8 disk drives being characterized as lightly-loaded or under-utilized so that their performance cannot be optimized. Thus, since run-time I/O demand may not be ideal, during system design or configuration it may be desirable to assume a value of Skew for a given storage device 110 to a number greater than about 1 as a contingency against times when demand may exceed a skew factor of about 1 for the given storage device 110, *e.g.*, Skew may be set to equal 1.4 (or 140%) such that the anticipated maximum number of